

PL-TR-96-2205

IMPLEMENTATION OF AN AUTOMATED METHOD FOR REGISTRATION OF FLIR IMAGES

S. M. Ayer

Radex, Inc.
Three Preston Court
Bedford, MA 01730

July 23, 1996

Scientific Report No. 1

Approved for public release; distribution unlimited

19961028 038



PHILLIPS LABORATORY
Directorate of Geophysics
AIR FORCE MATERIEL COMMAND
HANSCOM AIR FORCE BASE, MA 01731-3010

DTIC QUALITY INSPECTED 1

"This technical report has been reviewed and is approved for publication"

Edward C. Robinson
EDWARD C. ROBINSON
Contract Manager
Data Analysis Division

Robert E. McInerney
ROBERT E. MCINERNEY, Director
Data Analysis Division


This report has been reviewed by the ESD Public Affairs Office (PA) and is releasable to the National Technical Information Service (NTIS).

Qualified requestors may obtain additional copies from the Defense Technical Information Center. All others should apply to the National Technical Information Service.

If your address has changed, or if you wish to be removed from the mailing list, or if the addressee is no longer employed by your organization, please notify PL/IM, 29 Randolph Road, Hanscom AFB, MA 01731-3010. This will assist us in maintaining a current mailing list.

Do not return copies of this report unless contractual obligations or notices on a specific document requires that it be returned.

DISCLAIMER NOTICE



**THIS DOCUMENT IS BEST
QUALITY AVAILABLE. THE COPY
FURNISHED TO DTIC CONTAINED
A SIGNIFICANT NUMBER OF
COLOR PAGES WHICH DO NOT
REPRODUCE LEGIBLY ON BLACK
AND WHITE MICROFICHE.**

REPORT DOCUMENTATION PAGE

*Form Approved
OMB No. 0704-0188*

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)			2. REPORT DATE 23 July 1996		3. REPORT TYPE AND DATES COVERED Scientific Report No. 1		
4. TITLE AND SUBTITLE Implementation of an Automated Method for Registration of FLIR Images			5. FUNDING NUMBERS PE 62101F PR 7659 TA GY WU AG Contract F19628-95-C-0106				
6. AUTHOR(S) S. M. Ayer							
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) RADEX, Inc. Three Preston Court Bedford, MA 01730			8. PERFORMING ORGANIZATION REPORT NUMBER RXR-96072				
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) Phillips Laboratory 29 Randolph Road Hanscom AFB, MA 01731-3010 Contract Manager: Edward C. Robinson/GPD			10. SPONSORING/MONITORING AGENCY REPORT NUMBER PL-TR-96-2205				
11. SUPPLEMENTARY NOTES							
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release Distribution Unlimited			12b. DISTRIBUTION CODE				
13. ABSTRACT (Maximum 200 words) The implementation of a method for registration of digital images recorded by a 8-12 micron Forward-Looking Infrared (FLIR) camera is described. The software runs independently of commercial software packages under the Silicon Graphics IRIX 5.3 environment that is used to create and store the images as movies. It is capable of registering and drift compensating entire movies -- filtering against any number of reference images -- in one step in less than one second per frame per reference image, each of which is approximately 10,000 pixels. The package was written in C, and is herein described in detail from both methodology and implementation perspectives.							
14. SUBJECT TERMS Model validation, Infrared imaging, Image registration, Forward looking infrared (FLIR), Image processing, Weather					15. NUMBER OF PAGES 24		
					16. PRICE CODE		
17. SECURITY CLASSIFICATION OF REPORT Unclassified		18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified		19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified		20. LIMITATION OF ABSTRACT Unlimited	

TABLE OF CONTENTS

1. INTRODUCTION	1
2. THE PROBLEM	7
3. THE METHOD	9
4. RESULTS	11
REFERENCES	15
APPENDIX A. IMAGING HARDWARE	16
APPENDIX B. SOFTWARE IMPLEMENTATION	17

LIST OF FIGURES

<u>Figure</u>	<u>Page</u>
1. A high-contrast image of the test-target scene, made during daylight under clear skies	2
2. The low angle of the sun in winter causes flare across parts of the image, seen here on the calibration target beneath the transient plate	3
3. The target scene in the rain during daylight	4
4. The target scene during daylight under clear skies after a heavy snowfall	5
5. A natural landscape scene composed mostly of wooded areas, including hills in the background	6
6. A good example of banding due to FLIR sensor defects	8
7. The full-frame color classification of the test-target scene	10
8. Cropped sub-areas of the test-target cartoon	12
9. A column slice of the test-target scene vs. time, before registration	13
10. The same column slice vs. time as Figure 9, after registration	14
11. Algorithm Flow for Multi-Subarea Correlation	18

1. INTRODUCTION

A large database of Forward-Looking Infrared (FLIR) images exists for validation of the Modeling and Visualization software of the PL/GPAA Weather Impact Decision Aids (WIDA) Laboratory's Air Combat Targeting/Electro-Optical Simulation (ACT-EOS) project. The software's modeling process encompasses a complex series of computations of radiance for various features in a scene. Weather data is used to compute temperatures of features such as grass, asphalt, trees, and various targets of known composition, under given environmental conditions. Using known material descriptions, radiances are computed for features in the scene. Using known geometry of these features, 3D perspective scenes are rendered.

The FLIR experiments contain actual temperature data for the same type of objects under the same general set of environmental conditions that the ACT-EOS software models. Through comparison by feature of temperature vs. time curves under different environmental conditions, as shown in Figures 1 through 4, the accuracy of the model can be validated.

The experiment images are stored as JPEG-compressed movies created on a Silicon Graphics (SGI) workstation. The ability to analyze the images depends upon compensating for the FLIR's camera drift, amidst image-processing obstacles such as low signal-to-noise ratio and image distortion, as well as FLIR sensor noise compounded by compression and interlacing artifacts.

The primary goal was to correctly register the large body of movies in a short period of time with a minimum of human interaction, so that statistical analysis of the brightness of specific objects in the scene, from which temperature is easily derived, could be performed. The best of earlier schemes to register the FLIR movies generated unsatisfactory results with all but the cleanest of data-sets and consumed over 5 seconds per frame processed. Furthermore, it depended entirely upon a commercial visualization package.

The second goal was to develop an implementation that adapted automatically to different scenes. Originally, the test scenes were of a target site that consisted of relatively high-contrast (via cyclical heating) geometrically-shaped objects. Ultimately, registration and analysis must be performed with natural landscape scenes containing irregularly-defined features, such as vegetation and roads, as shown in Figure 5.



Figure 1. A high-contrast image of the test-target scene, made during daylight under clear skies. The calibration target (upper left, with legs) contains three aluminum plates. From top to bottom, they are ambient temperature, heated cyclically, and heated intermittently. The ground plates (foreground, on cubes) are near ambient temperature.




Figure 2. The low angle of the sun in winter causes flare across parts of the image, seen here on the calibration target beneath the transient plate.




Figure 3. The target scene in the rain during daylight. Rain has cooled every feature in the image to nearly the same temperature, reducing the contrast to almost zero.




Figure 4. The target scene during daylight under clear skies after a heavy snowfall. Most of the ground plates are obscured by snow. Note the warm shadows at the bases of the objects where the snow is less deep.




Figure 5. A natural landscape scene composed mostly of wooded areas, including hills in the background. Cultural features in the cool (dark) clearing in the foreground include a small building (lower right), a road, and a number of utility poles.

2. THE PROBLEM

The need for a registration scheme arose from the FLIR's tendency to drift vertically, which makes analysis of the experiment data practically impossible unless performed manually on a frame-by-frame basis. Attempts at tracking a horizontal edge in order to detect and correct the drift revealed that the FLIR also drifts horizontally. Further investigations showed that the FLIR images suffer from blurry edges.

Due to mismatched gains from the two rows of the FLIR sensors, the FLIR images also appear to be striped. Interlacing enhances this effect, making the stripes two pixel-rows in height. Compressing the images compounds the problem by causing any given dark/dark or light/light pair in a stripe to have identical row sums, making it virtually impossible for the software to "distinguish" individual scan-lines when correlating.

Furthermore, a large number of the images in the data-set also have a column offset of one or two pixels because of a sensor time-delay error, giving edges a saw-toothed appearance, as in Figure 6.

Image registration requires a modicum of clarity in the reference image. Most methods assume that a noise-free reference is unattainable [Eversole, Nasberg, 1983] and that rectification must be performed on an image prior to registration to correct for different perspectives and resolutions, typically the result of a moving camera, e.g. satellite image processing.

Rectification of the FLIR images is not necessary since all of the images come from the same fixed-position camera. However, a great deal of edge-enhancement and noise-filtering is required to produce a reference image. Unfortunately, even extensive pre-processing of a FLIR image results in a poor-quality reference that can result in numerous false fixes. A useful rule adopted during development was that if a human could not determine the exact position of an edge in a reference image, then the software would certainly have difficulty making the same determination.




Figure 6. A good example of banding due to FLIR sensor defects. Each stripe is two pixels in height. Note jagged vertical edges.

3. THE METHOD

The basic image registration method employs correlation of a reference image to a filter image [Pratt, 1978] using Fast Fourier Transforms [Schowengerdt, 1983]. Both reference and filter images were used full-frame (640x480 pixels). The reference image is reversed and normalized, and an FFT is performed. Each frame of the movie is used as a filter. An FFT is performed on the filter image, the product of the two is calculated, and then an inverse FFT is run on the result. The maximum of this function represents the registration of the two images, or the number of pixels that the filter image drifted from the canonical position of the reference position.

Initial efforts to register the movies used a “clean,” high-contrast frame directly from the movie being processed as a reference image. This method produced mixed results due to the image’s inherent sensor, interlacing, and compression noise. Thus, success in registering images using a canonical FLIR image as a reference depended heavily upon the filter image’s similarity in contrast, signal-to-noise, and banding orientation.

To overcome this limitation, two reference images were constructed from the original FLIR references -- one for each target scene in the data-set -- with a commercial off-the-shelf Geographic Information System (GIS) package, ARC/INFO. The scene was classified by hand, with each feature-class mapped to a different color, e.g. grass was mapped to green, target legs were mapped to blue, and the calibration plates were each mapped to a separate shade of grey, as shown in Figure 7.

This noise-free reference provided sharp edges and eliminated the difficulty of choosing a single good reference from the FLIR images. It made the registration method much stronger by allowing individual correlation of each feature-class in the image.

This method produced accurate results even with poor quality filter images, but ran slowly -- approximately one minute per frame -- because of the heavy computational burden of correlating (two 640x480 images) for each feature (the target reference has 11 features, the landscape reference has 20 features).

In order to increase speed, the colored reference image, hereafter referred to as the “cartoon,” was cropped to include only features that are nearly always visible above the noise in the FLIR images. Each filter image was cropped accordingly for the registration. A rectangular cropped area of roughly 10,000 pixels -- still including roughly 60% of the original reference features -- proved sufficient to successfully register movies in which the cropped portion was distinguishable in every frame.

The single cropped-cartoon method was not robust enough to handle the range of environmental conditions in the ACT-EOS data-set. This data-set includes experiments made over a four season period in which changing foliage, solar angle, and weather at times obscure or blur parts of the




Figure 7. The full-frame color classification of the test-target scene.

cropped area. Since operating with a smaller reference image provided a 40-fold performance gain, it pointed toward the use of multiple cropped-reference-images, with a separate registration for each. Figure 8 shows three sample crops of the original cartoon.

The multiple sub-area correlation [*Lo, Gerson, 1979*] variation provides quick registration with the ability to avoid false fixes by feeding the multiple registration data to a decision-making function that essentially determines which of the data is least aberrant, taking advantage of the fact that the drift in the FLIR movies is smooth on a frame-to-frame basis.

4. RESULTS

Performance data were obtained with an executable compiled and run on a SGI Indigo 2 with a 200MHz MIPS R4400 CPU and 128MB RAM, running IRIX 5.3.

Sample runs were made with three cartoons, with a slight modification that forced the software to run all three possible registrations for each frame. Cartoon sizes and complexities were: 80x150 pixels, 3 colors/features; 155x63 pixels, 13 colors/features; and 55x110 pixels, 4 colors features. Typical elapsed time (measured using the SGI prof utility) for registration of an 863 frame movie is 1400 seconds, or 1.6 sec/frame.

The registration software is effective to within 4 pixels. Figures 9 and 10 show before- and after-registration column slice vs. time images for one experiment, respectively. Overall vertical drift in the uncorrected movie is 18 pixels. The corrected movie drifts a total of 4 pixels, despite visible evidence of image size distortion over the course of the movie, and the signal-to-noise varies widely over the course of the movie.

Given this performance, the goal of correctly registering FLIR images quickly with little human interaction was achieved. The software registers -- within the aforementioned tolerance -- a 900-frame movie, producing a new set of drift-corrected frames, in roughly 20 minutes. The only setup required is entering the movie name and the desired cartoon reference(s) in a configuration file.

The goal of automatic adaptation to different scenes was not achieved, due to the fact that a new cartoon reference must be hand-made for each new scene introduced to the registration software. Once the cartoon is made, this reference must be cropped manually into a sufficient number of sub-images to properly register the scene. "Sufficient number" may need to be defined experimentally. Once past this process, though, the software itself registers the new scene automatically.




Figure 8. Cropped sub-areas of the test-target cartoon. Crops were made to maximize the registerable-feature to area ratio. Note calibration target legs without plates (middle above). This sub-area avoids sun flare on the plates during winter scenes.




Figure 9. A column slice of the test-target scene vs. time, before registration. The column selected bisects the calibration plates. Vertical motion represents camera drift.




Figure 10. The same column slice vs. time as Figure 9, after registration.

REFERENCES

Pratt, W. K., "Digital Image Processing", New York, John Wiley and Sons, 1978, 750 pp.

Showengerdt, R. A., "Techniques for Image Processing and Classification in Remote Sensing", Orlando, Academic Press, pp. 249, 1983.

Lo, T. K. and G. Gerson, "Guidance System Position Updated by Multiple Subarea Correlation", SPIE Proceedings, Volume 186, May, 1979.

Eversole, W. L. and R. E. Nasberg, "Maximum Likelihood Estimation for Image Registration", SPIE Proceedings, Volume 432, August, 1983.

APPENDIX A. IMAGING HARDWARE

The camera used to make the experiment images is a FLIR Systems FLIR 2000, which produces NTSC video images in the 8 to 12 micron IR spectrum. These images are recorded to S-VHS tape and digitized as 640x480 pixel SGI rgb files -- by Silicon Graphics's Indy Video -- at a frame-rate of one per minute. The movies are created from the image files by the SGI utility makemovie.

APPENDIX B. SOFTWARE IMPLEMENTATION

The software implementation was written entirely in ANSI C. SGI library calls were used for movie frame extraction and restoration, and all FFT operations. Library dependencies are all outlined in the Makefile. All libraries and shared objects are part of the SGI development option release except for libmovieutils.a, which was built for this application and resides in the source directory. Five C source files and three header files are necessary. On an SGI workstation with the X, Motif and dmedia packages, building the executable should be as simple as running make in the source directory.

The software takes one command-line argument, the name of a configuration file that contains one movie filename, two integer drift tolerances (see explanation below) and a variable number of reference image filenames, respectively. Comment lines start with '#'. The only format restrictions are that 1) the data appear in this order, 2) that no comments follow the reference image filenames; these must be the last entries in the file, and 3) that only one item is placed on a single line.

The movie file must be an SGI movie, and the reference images must be SGI rgb files. See Figure 11 for global algorithm flow diagram.

The first major step in the implementation deals with pre-processing the reference images. Since the number of reference images is varied depending upon the "difficulty" of the experiment, i.e. quality of the image and weather conditions, config_data() builds a linked list of data structures (struct canon_data) that contain all pertinent information on the cartoon. Each cartoon is processed by call to prepare_canonical().


Prepare_canonical() assumes that the cartoons have been cropped with a unique, monochrome border at least one pixel deep. The area inside the border will be referred to as the "cartoon window."

Each node in the linked-list contains the overall dimensions of the cartoon window, and the location of the cartoon window, in x and y coordinates of the bottom left corner of the cartoon window.
NOTE: In SGI rgbs, the y-coordinate increases from the bottom of the image, opposite the standard X-window convention. All coordinates are treated in this fashion.

All image representations are "flattened" into one-dimensional arrays in the implementation so that they are consistent with the SGI movie frame extraction utilities, which return 1D arrays of unsigned chars.

Once the cropped area is identified, prepare_canonical() builds an array of structs (struct bucket) to be attached to the canon_data node for that cartoon. Each element in this array contains a color's red, green, and blue values, the total number of pixels of that color in the cartoon window, and a "bit bucket" array. This is a binary representation -- literally 1 if a pixel is that color, 0 otherwise -- of the feature. Once the array is reversed and normalized, an FFT is performed in place. The first half of the correlation process is then complete.

Algorithm Flow For Multi–Subarea Correlation



For each sub-area cartoon of the reference image, the filter may be cropped and a correlation performed. The decision function validates the registration results, then either calls for another correlation with a different sub-area, or uses the registration value to drift-correct the frame.

Figure 11. Algorithm flow for multi-subarea correlation.

The second major step in the registration software is the acquisition of individual movie frames, which are the filter images in the correlation. `Movie_to_images()` opens the movie, and each frame is extracted, de-compressed and de-interlaced as necessary by calls to SGI library functions. It is then stripped to a single greyscale value per pixel array.

Both elements of the frame's correlation are then in place, and the registration process for the frame is performed. First, an FFT is performed on the filter image array. Next, a correlation is done with each feature's pre-processed (normalized, reversed and FFTed) bit-bucket array.

The results of these identical operations are combined by adding "across" each feature's individual correlation. A new array is built containing the sums of the raw correlation values at that pixel, e.g. element 0 of the sums array contains the correlation value in element 0 of each feature's correlation result.

The index of the peak value in this array is the registration value for the frame.

This drift value is broken into x and y components for examination by `validate_drift()`. This function uses the drift tolerances in the configuration file to check each frame's perceived drift -- in pixels -- for plausibility in order to catch possible misregistration.

The first of the two drift tolerance values is maximum absolute drift allowable in either the x or y direction. The second value is the maximum relative drift in either direction from frame-to-frame. The tolerances can be set most effectively if the user has a modicum of familiarity with the nature of drift -- both the magnitude and smoothness -- in the movie being registered.

Any calculated drift that violates either tolerance check causes `validate_drift()` to "doubt" the veracity of the registration with the current cartoon, and a recursive call is made to re-register the frame with the next cartoon in the linked-list. The implicit assumption is that the current cartoon registers a part of the scene that is occluded or blurred in the frame, so registering with another cartoon, in a different area of the image, should be tried until a "believable" answer is returned. If every registration is perceived to have failed, i.e. all cartoons have been exhausted, then the image is assumed to be completely obscured by noise and unregisterable. In this case, the frame is held in place relative to the previous frame by using the same amount of drift compensation.

Once the frame is registered, the software shifts the image the determined amount. Pixels at the ends of the 1D representation are "wrapped," e.g. if an image is shifted six pixels to the right, the last six pixels in the image (top right) become the first six pixels in the corrected image, or bottom left.

The new, drift-corrected frame is written out as a new SGI rgb image. New frames are placed in a directory `registered-images` under the user's current working directory. The name is based upon the movie with ".reg" and ".xxxx" inserted, where xxxx is a four-digit frame number, e.g. `flir.mv` produces corrected frames `flir.reg.0000.rgb`.

A registered, drift-corrected frame is produced for each frame in the original movie.

Two tools were also developed to check the software's performance.

One tool was composed by dividing the functionality of the registration software. The first executable performed the registration for all available cartoons without validating the data, instead placing the raw row/col drift values into a file. The second half performs the actual drift compensation on each frame of a movie using output from the registration part of the tool. This was used to evaluate the registration software by comparing raw data to actual frame movement when playing the unregistered movie, a somewhat subjective but highly effective process. Several shortcomings were discovered during the development phase using this technique.

The other tool provides a more graphic evaluation of the registration software. This tool opens each frame of a movie, slices the same column -- or row -- from the frame, constructing one new image. The output image is effectively a position vs. time graph. By selecting a slice with a bright feature, a graph of either vertical or horizontal drift can be made. This tool allowed for direct comparison between drift in unregistered and registered versions of the same movie.